

Virtex-II Pro SEE Test Methods and Results

David Petrick¹, Wesley Powell¹, James W. Howard Jr.², Kenneth A. LaBel¹

¹ NASA/Goddard Space Flight Center, Greenbelt, MD 20771

² Jackson & Tull, Chartered Engineers, Seabrook, MD 20706

Abstract

The objective of this coarse Single Event Effect (SEE) test is to determine the suitability of the commercial Virtex-II Pro family for use in spaceflight applications. To this end, this test is primarily intended to determine any Single Event Latchup (SEL) susceptibilities for these devices. Secondly, this test is intended to measure the level of Single Event Upset (SEU) susceptibilities and in a general sense where they occur.

The coarse SEE test was performed on a commercial XC2VP7 device, a relatively small single processor version of the Virtex-II Pro. As the XC2VP7 shares the same functional block design and fabrication process with the larger Virtex-II Pro devices, the results of this test should also be applicable to the larger devices. The XC2VP7 device was tested on a commercial Virtex-II Pro development board. The testing was performed at the Cyclotron laboratories at Texas A&M and Michigan State Universities using ions of varying energy levels and fluences.

I. INTRODUCTION

The Xilinx Virtex-II Pro is a SRAM-based platform FPGA that embeds multiple microprocessors within the fabric. The variety and quantity of resources provided by this family of devices make them very attractive for spaceflight applications. However, these devices will be susceptible to SEEs, which must be understood in order to be mitigated.

To use the Virtex-II Pro reliably in space applications, these devices must first be tested to determine if they are susceptible to SEL, the degree to which they are susceptible to SEU and single event transients (SET), and how these effects are manifested in the device. With this information, mitigations schemes can be developed and tested that address the specific susceptibilities of these devices.

Due to the complexity of these devices, it is advantageous to partition the SEE testing into a series of tests, with each test focusing on a specific set of functions. This paper presents the objectives, approach, and results for a first-order

SEE test for the Virtex-II Pro. The first objective of this test is to determine if the devices are susceptible to SEL, and therefore unsuitable for space applications. Secondly, the coarse SEE test has the objective of determining the general SEU/SET performance of the devices, although with limited insight into the exact upset mechanisms.

The coarse SEE test uses a commercial-off-the-shelf (COTS) Virtex-II Pro evaluation board provided by Memec, with a single processor XC2VP7 FPGA. The FPGA on this board is replaced with a delidded device and partially covered with a shield. During SEU testing, this shield is placed on the device to only expose certain portions of the logic, routing, configuration memory resources, MGTs, or PowerPC.

The logic design consists of a combination of an existing test application, the Xilinx Bit Error Rate Test (BERT) application, along with a set of simple test structures. The BERT application uses the embedded processor along with a pair of RocketIOTM transceivers to perform a bit error rate self test. Using this application, upsets in the transceivers will be manifested as bit or link errors, while processor upsets are manifested in resets, communication, and display errors. The test structures are then included to test the configurable logic blocks (CLBs), embedded multipliers, and BlockRAM. Here an identical set of test structures, one in a shielded area and one in an exposed area, are driven with an identical pseudo-random data pattern. The outputs of these test structures are then compared using circuitry in a shielded area. Upsets are then manifested as mismatches.

II. TEST DETAILS

A. Virtex-II Pro Device Properties

The FPGA used will be the commercial Xilinx Virtex-II Pro XC2VP7-6FG456C device. This device includes a single embedded PowerPC processor, 4.4 million configuration bits, 792 kB of BlockRAM, 8 RocketIOTM Multi-Gigabit Transceivers (MGT), 4 Digital Clock Managers (DCM), and 44 dedicated 18x18

multipliers [1]. The package used will be the wire-bond 456-pin ball grid array (FG456).

In order to prepare the part for radiation testing, the die needs to be exposed through an acid etching process. Analytical Solutions located in Albuquerque, New Mexico, successfully performed the etching process using sulfuric acid at high temperature. An acid etched Virtex-II Pro FPGA is shown in Figure 1.




Figure 1 – Acid Etched Device

B. Memec Development Board

The development board used for these test is the Memec DS-KIT-2VP7FG456 (Figure 2), which contains one soldered FPGA along with external RAM, PROMs, RS-232 port, JTAG connectors, MGT drivers and connectors, oscillators, power converters, and various user switches. Also included with this board is a prototyping daughter card (DS-KIT-P160-PROTO), which was populated with RS-422 line drivers to generate discrete pulses that indicate detected upsets.

Since this board was not designed for tests of this sort, some minor modifications were necessary. The user needs to have control over the program, reset, and mode pins without needing to enter the radiation chamber. These switches were removed and replaced with headers for fly-lead wires leading to a switch box, which can be extended to the user area at the test facility. Next, due to an impedance load problem, the existing DC/DC power converters

were electrically removed from the board. This enables the board to be power by an external power supply. Finally, the FPGA that comes pre-installed on the board was replaced with an acid-etched-delidded FPGA. The FPGA replacement was performed at the Goddard Space Flight Center by the Parts, Packaging and Assembly branch.

Three boards were populated with delidded FPGAs. During radiation testing, all boards were used to duplicate upset mechanisms in order to give the collected data more statistical significance and an understanding of possible device-to-device variation.




Figure 2 – Memec Development Board

C. FPGA Design

The FPGA circuitry that undergoes SEE testing included the following Virtex-II Pro functional elements: (a) the PowerPC processor, (b) MGTs, (c) BlockRAM, (d) dedicated multipliers, (e) CLBs, and (f) configuration RAM. These functional elements were tested using the combination of the BERT reference application and standalone test structures. The BERT reference application, which was modified to accommodate this board and test, tested the operation of the processor and the MGTs. For testing purposes, the MGT cables were hooked up in loopback (i.e. TX -> RX). This allows the transmitted pseudo-random data

to be compared at the receiver, and detect any bit errors.

The BERT application consists of: (a) a FPGA image containing processor peripherals and MGT support circuitry, (b) embedded software running on the PowerPC processor, and (c) user interface software running on a standalone PC. This PC communicates with the Memec board via a RS-232 interface. Using this application, MGT upset events are observed as bit errors or link failures, which are displayed on the user interface log window. Processor errors are detected as software malfunctions, also indicated on the user interface log window (i.e., corrupted RS-232 communications, software hanging, etc.).

The BlockRAM, multipliers, configuration RAM, and CLBs were tested with a dedicated test structure. This test structure, which is depicted in Figure 3, consists of dual sets of circuitry, with one set in the exposed area and the other in the shielded area. Both are driven with a pseudo-random data generator. The outputs of these sets of circuitry are then compared. If detected, the comparators generate error pulses indicating upsets in the exposed circuitry.

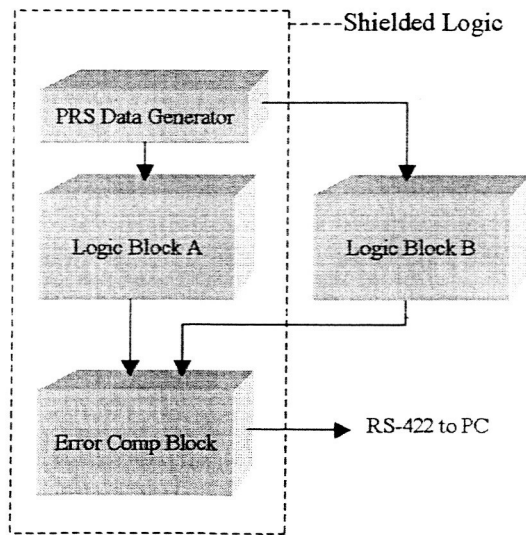


Figure 3 – SEU Comparison Logic

This test structure can either be integrated with the BERT application, resulting in a single FPGA image, or programmed as a standalone function. During testing, this will enable the flexibility to isolate different functionality of the FPGA, or run all logic at once. By floorplanning the designs accordingly, different parts of the

device/design can be shielded to suppress any unwanted SEEs.

Two onboard oscillators are used by the FPGA to derive the internal clock frequencies using DCMs. A 100 MHz oscillator is connected to one DCM that uses this as a reference to supply the PowerPC with a 200 MHz clock and the FPGA fabric with a 50 MHz clock. The other DCM is connected to a 125 MHz oscillator, which controls the MGT reference clock that derives the data rate. The MGTs were set to run at 2.5 Gbps.

D. Experiment Setup and Instrumentation

As illustrated in Figure 4, the test instrumentation consists of a pair of Windows-based computers, a HP6624A quadruple-output power supply, and a switch box. Through a custom graphical user interface (GUI), the instrumentation computer allows the user to control the power supply settings [2]. It also strip charts the voltage and current for each output and maintains a count of discrete pulses, which are connected to error signals from the device under test (DUT) [2].

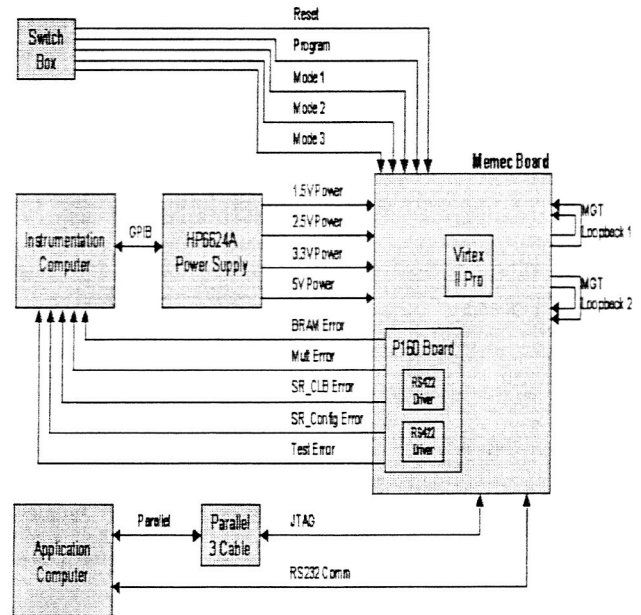


Figure 4 – Instrumentation

The user has the option of programming the FPGA from the PROMs or via JTAG. When programming via JTAG, the application computer controls the programming and verification of the FPGA designs through the Xilinx IMPACT software [1], which controls the Parallel 4 programming cable, also a product of

Xilinx. This computer also runs Tera Term Pro, which is used as the user interface software for the BERT application and records log files pertaining to each test.

Both the instrumentation and application computers are located in the test chamber to ensure signal integrity with shorter cables. A third computer, not shown in Figure 4, is used to control these computers through a virtual network connection. With the switch box and control computer located in the user area, the user has full control of this test setup.

E. Device Shielding Techniques

As mentioned before, the device can be covered with a shield to leave only the areas of interest exposed to the radiation. Two L-shaped, 3mm brass shields were milled for these tests. When the shields are properly configured any target area on the FPGA can be left exposed. The shields are held in place with Kapton tape.

Using a high-power stereomicroscope, very precise placement can be achieved. Since the die is exposed, some of the Virtex-II Pro logic resources are visible. In particular, tight circuit blocks on the outer edge of the die represent the DCMs and MGTs. On the other hand, the most important Virtex-II Pro resource, the PowerPC, is not visible due to the fact that it is embedded in the FPGA. Hence, this poses a problem when trying to expose, or shield it with the masks.

The methodology employed for mask placement was based on strategic floorplanning of the design. The FPGA design was partitioned with the Xilinx Floorplanner tool. As shown in Figure 5, the BERT application logic and the majority of the SEU test structure are isolated on the left half of the FPGA. Other logic functions, such as the UART, system reset, clock startup, and JTAG logic are pushed to the bottom of the FPGA. The part of the SEU test structure to be exposed to radiation, 'Logic Block B' (as discussed in *C. FPGA Design*), is placed in the upper right-hand section of the FPGA.

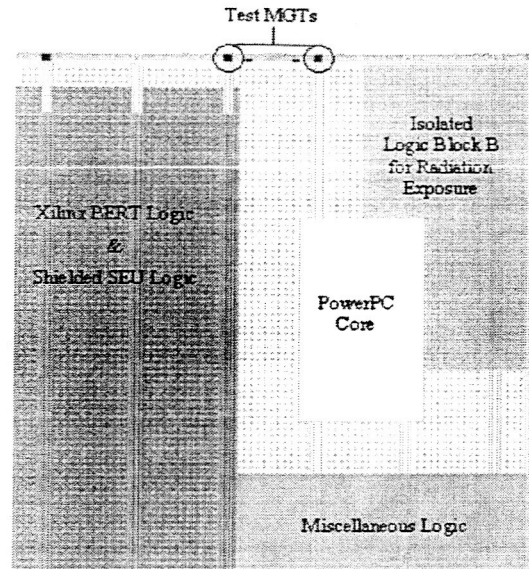


Figure 5 – FPGA Design Floorplan

Two main shielding configurations were used during the coarse SEE testing. The first was when testing for PowerPC SEEs and test structure SEUs. One L-shaped mask was placed to shield the left half and bottom part of the FPGA, leaving the PowerPC and 'Logic Block B' exposed. A "ball-park" placement of the mask, depicted in Figure 6, was used since the PowerPC Core is not visible. This was the reasoning for the format of the design floorplan. A small buffer of unused logic surrounds the PowerPC in order to allow for a small amount of error in the mask placement.




Figure 6 – PowerPC Mask Placement

The second shield configuration, depicted in Figure 7, was used when testing for MGT SEEs. Both shields were used in order to mask off only

one MGT. With the aid of a good microscope, the placement of this mask configuration is very accurate since the MGTs are visibly confined to a small circuit block.




Figure 7 – MGT Mask Placement

III. TEST RESULTS

A. SEL Testing

The main goal of the coarse SEE test for the Virtex-II Pro was to determine if the device will enter a latch-up state under radiation conditions. When testing at the Cyclotron Laboratories at Texas A&M or Michigan State Universities, no destructive SEL event was observed to a LET of 53.9 MeV-cm²/mg and a fluence of 10⁷ ions/cm².

During SEL testing, some interesting observations were made. While being irradiated, the internal current (I_{CCINT}) slowly rises. Once a current of ~3.4 A is reached, the current drops to 0 A, jumps back up to nominal, and then continues to ramp. It was determined that the device was reloading the configuration file from the on-board PROMs. The over-current protection setting for the power supply was set to 5 A and the power supply was capable of sourcing the 3.4 A without any voltage sag. So this characteristic is not a result of the power supply clipping the current or voltage sagging at the device.

As shown in Figure 8, the current cycling is symmetric and the voltage does not sag. Through additional tests, it was found that the rate at which the current ramps was a function of the radiation conditions, the size of logic design loaded, and the amount of die exposed. Also, this event does not occur if the FPGA is powered without an initial configuration.




Figure 8 – Current Ramping Characteristic

Also during SEL testing, a configuration readback was performed after each run to determine the number of upset configuration bits. This was simply performed by clicking the 'Verify' command in the Xilinx ISE 6.1i iMPACT tool [3]. This command counts the number of differences found in the configuration data. Figure 9 shows the results of this data. Note that there are ~4.4 million configuration bits for this particular part.

To make an attempt to account for the MGT configuration bits, a mask was placed on the FPGA to let only four MGTs exposed. The cross-section is reduced by approximately a factor of 10. This agrees with the fact that about 90% of the die was shielded from the radiation.




Figure 9 – Configuration Memory Bit-Error Cross-Section

B. SEFI Testing

The purpose of this testing phase was to document any observed event that would be classified as a Single Event Functional Interrupt (SEFI) such as a failure in a MGT link or the PowerPC log data getting corrupted, skipping instructions, or halting. The main objective was to focus on the PowerPC operation. For this case, the mask configuration showed in Figure 6 was used. Although the exact location of the PowerPC resources are unknown, this mask placement served as a rough estimate for initial testing purposes.

In Figure 10, the red curve (No PPC) represents the SEFI data collected when the PowerPC was completely shielded from the radiation. The blue curve (With PPC) represents the SEFI data observed when the PowerPC was exposed to the radiation. When testing with the PowerPC core exposed, the flux of the beam was turned down very low, on the order of $2.5E2$ ions/cm²/sec. Initially the flux was set to approximately $3.2E5$ ions/cm²/sec, but once the beam was turned on, the log file immediately halted. Therefore, in order to collect statistically significant SEFI data, the flux had to be decreased.




Figure 10 – SEFI Cross-Section

C. MGT SEU Testing

The goal of this test was to gather upset data on the MGTs. The data pattern used to drive the MGT transmit ports was a pseudo-random pattern of $1+X^6+X^7$. The MGTs were running at a data rate of 2.5 Gbps. In order to isolate one MGT, specifically MGT6 on the XC2VP7, the mask configuration showed in Figure 7 was used throughout testing.

For each run, the number of MGT bit errors was recorded. This data was extracted from the PowerPC log file. The run was terminated upon MGT link failure. Figure 11 shows the cross section data collected with a Weibull fit to that data. The threshold LET was found to be about $0.1 \text{ MeV-cm}^2/\text{mg}$ and the saturation cross section was approximately $2.6 \times 10^{-5} \text{ cm}^2$. The complete parameters of the Weibull fit are shown in the figure legend.




Figure 11 – MGT Bit-Error Cross-Section

IV. DISCUSSION

The coarse SEE test was successful and most objectives were met. The only data that was not collected was related to the logic test structure designed to detect SEUs in the BlockRAMs, MULT18x18s, CLBs, and configuration RAM. During a run, once the beam was turned on, all counters eventually started incrementing at a significant pace (more than should be expected for the gate count involved). Since the counters were rolling over the max setting of 65536, it was also impossible to calculate the actual amount of errors per run. In this type of test setup, the SEU error counters should only increment at a slow rate (i.e., 1 error every few seconds). During some tests, one of the counters would halt for a few seconds, but then start up again.

There are a few explanations for the counters to increment in this fashion. First, the FPGA designs used were not implemented with triple modular redundancy (TMR), nor is the configuration memory being dynamically reconfigured, or scrubbed. Therefore, an accumulation of configuration upsets could lead

to this type of event. Second, it is possible that the Input-Output Blocks (IOB) controlling the error counters may have been hit [4]. Third, if the DCM was hit, hence taking out the clock, the error lines could possibly have been stuck high. But in the latter two cases, it does not seem statistically possible due to the size of an IOB or DCM with respect to the entire FPGA and the frequency with which these types of events occurred.

The constant current ramping characteristic discussed in *SEL Testing* is most likely caused by the configuration bits gradually turning on at a constant rate. However, it is statistically abnormal that the configuration bits all turn "on", resulting in the monotonic data collected. It seems that some of the configuration bits should turn "off" when exposed to radiation, resulting in a slight drop in current.

The second event observed in the current ramping data is the reconfiguration of the device once a current of 3.4 A is reached. This device is rated to 8 A, so the FPGA is not reaching its physical limit [1]. Post-testing and analysis has shown that the remote sensing is not completely accurate due to modifications made to an existing setup to accommodate the Memec board [2]. The power cable used in these tests consists of a 40-pin ribbon cable, which is wired for remote sensing feedback, and custom 12-inch fly-leads to connect to the Memec board. The remote sensing is done at the end of the ribbon cable, or before the fly-lead extension. Therefore, the power supply will keep the voltage at the end of the ribbon cable constant at the programmed value. This is why the plot in Figure 8 shows V_{CCINT} remaining constant at 1.5 volts. However, at high currents, due to a small IR drop across the fly-lead extension, the core voltage of the FPGA is falling below the minimum voltage required to properly power the device. This causes a power-on reset (POR) sequence to occur in the FPGA, which explains the reconfiguration [1, 3]. The impedance of the fly-lead extension cable was measured at 66 m Ω . So at a current of 3.4 A, there is a 230 mV voltage drop across the fly-lead cable. The POR is occurring when V_{CCINT} drops to 1.27 V (15% drop). However, when testing at a nominal current of 500 mA, there is only a 33 mV voltage drop, keeping the FPGA voltage well above the $\pm 5\%$ recommended operating condition [1].

After each test run, the configuration bit errors were counted and the action required to reestablish the functionality of the device was documented. During SEL testing, over 400,000

configuration bit errors ($>10\%$) were recorded twice and the JTAG link failed twice. Both types of occurrences are probably due to configuration errors in the JTAG circuitry. The results on reestablishing device functionality are as follows: reset PowerPC – 28%, toggle Prog pin – 70%, cycle power – 2%. Most of the software resets were successful when the PowerPC was shielded during MGT SEU testing and cycling the power was necessary when performing latch-up tests.

V. CONCLUSION

The commercial-grade Virtex-II Pro did not enter a latch-up state during these tests. However, a preliminary conclusion is that the MGTs and embedded PowerPC have a high susceptibility to the heavy ion radiation since SEFIs occurred too quickly to collect enough substantial data. Due to the limitations on isolating the PowerPC with a shield configuration, new test methods need to be developed in order to gather more conclusive data on its operation.

The coarse SEE testing phase also helped expose other areas of the test procedure and setup that need adjustments. In particular, the remote sensing issue will be fixed for future tests. Aside from that, the largest problem encountered was the inability to continuously reconfigure the configuration memory in order to eliminate any bit-flips. The Memec board was a good test bed for the coarse SEE test, however, in order to allow more flexibility for future tests, we will need to migrate to another board more suited for radiation testing. This will permit SEU monitoring, local (to the device) voltage and current sensing, TMR, and partial reconfiguration. Through TMR and scrubbing, accurate results can be attained indicating how the device will operate in a radiation environment [4, 5, 6].

VI. ACKNOWLEDGEMENTS

The authors would like to thank the following for their support on this project: (1) Tim Irwin of QSS Group Inc. for his time spent learning the process of delidding the Virtex-II Pro, placement of the masks, and test facility assistance (2) Rich Katz of GSFC and Igor Kleyner of Orbital Sciences Corporation for assisting with the instrumentation setup (3) NASA Electronic Parts and Packaging Program

(NEPP) for supplying the beam time at Michigan State University.

VII. REFERENCES

- [1] "Virtex-II Pro Platform FPGA Handbook", Xilinx, Oct. 2002
- [2] J. Wang, W. Wong, S. Wolday, B. Cronquist, J. McCollum, R. Katz, I. Kleyner, "Single Event Upset and Hardening in 0.15 μ m Anitfuse-Based Field Programmable Gate Array", MAPLD'02, Laurel MD, 10-12 Sept. 2002
- [3] C. Yui, G. Swift and C. Carmichael, "Single Event Upset Susceptibility Testing of the Xilinx Virtex-II FPGA", MAPLD'02, Laurel MD, 10-12 Sept. 2002
- [4] G. Swift, S. Rezgui, J. George, C. Carmichael, M. Napier, J. Maksymowicz, J. Moore, A. Lesea, R. Koga, T. Wrobel, "Dynamic Testing of Xilinx Virtex-II Field Programmable Gate Array's (FPGA's) Input Output Blocks (IOBs)", NSREC'04, 19-23 July 2004
- [5] C. Carmichael, "Correcting Single-Event Upsets through Virtex Partial Reconfiguration", Xilinx Application Note XAPP216, June, 2000
- [6] E. Fuller, M. Caffrey, A. Salazar, C. Carmichael, J. Fabula, "Radiation Characterization, and SEU Mitigation, of the Virtex FPGA for Space-Based Reconfigurable Computing", NSREC'00, Reno NV, 24-28 July 2000